arXiv:1511.06489v1 [cs.CV] 20 Nov 2015

A Simple Hierarchical Pooling Data Structure for Loop Closure

Xiaohan Fei Konstantine Tsotsos Stefano Soatto
UCLA Vision Lab, University of California, Los Angeles, CA 90095
{feixh, soatto, konstantine}@cs.ucla.edu

Abstract

We propose a data structure obtained by hierarchically averaging bag-of-word descriptors during a sequence of views that achieves average speedups in large-scale loop closure applications ranging from 4 *to* 20 *times on benchmark datasets. Although simple, the method works as well as sophisticated agglomerative schemes at a fraction of the cost with minimal loss of performance.*

1. Introduction

We tackle the problem of *loop closure* in vision-based navigation. This is a particular classification task whereby a training set of images is indexed by location, and given a test image one wants to query the database to decide whether the former is present in the latter, and if so return the indexed location. This is closely related to *scene recognition*, where the focus is on a particular instance, as opposed to an object class (we want to determine whether we are at particular intersection in a given city, not whether we are at *some* intersection of *some* urban area). As such, test images are only subject to *nuisance variability* due to viewpoint, illumination and partial occlusion from moving objects, but otherwise there is no *intrinsic* (intra-class) variability.

The state-of-the-art for image retrieval is based on convolutional neural network (CNN) architectures, trained to marginalize nuisance and intrinsic variability. While some of this variability is managed by design, for instance through the use of linear convolutions and pooling, the rest is *learned away* by means of large annotated datasets. Unfortunately, existing convolutional architectures are not well matched to the loop closure problem: In a discriminatively trained network, the compositionality property afforded by linear convolutions, while critical to model intra-class variability, is detrimental to loop closure, as there is no intrinsic variability. At the same time, a CNN does not respect the topology of data space at higher levels of the hierarchy, since filters at any given layer are supported on the entire feature map of the previous layer. In loop closure, locality is key, and while one could retrieve from the feature map the locations that correspond to active units, this requires some effort [30]. In this sense, CNNs are too much and too little for loop closure as they are not designed to specifically tackle the problem. It is not a coincidence that the state-of-the-art in loop closure consists of variants of bag-of-words (BoWs) and inverted indices [10, 22] that were ubiquitous in image retrieval before the advent of CNNs, which are powerful for object identification, not for applications like SLAM which we focus on in this paper. Unlike generic retrieval, however, BoWs have remained the most popular.

Given the critical importance of loop closure in location services ranging from smartphones to autonomous vehicles, we focus on its peculiarities, and attempt to harvest some of the components of neural networks to improve the state-of-the-art. Stripped of the linear convolutions (we do not need to model intrinsic variability) and ReLu (occlusions can be modeled as a subset selection, rather than a marginalization as suggested in [26]), what we have left is a *hierarchical spatially pooled data structure built upon local photometric descriptors [13, 21]*. There are no filters, and no learning other than the trivial pooling of local descriptors. Motivated by this intuition, we propose a new hierarchical representation for loop closure, detailed in Sec. 2.

Loop closure is also closely related to location, or "place," recognition [40, 38, 6] and large-scale visual search [25, 16, 5, 17], but with some important restrictions.

First, both previous data (training images) and current (test, or query) data are usually available as time-indexed sequences, even if they are captured by different agents, and training images may be aggregated into a "map" [18] or reduced to a collection of "keyframes" [24]. Second, as a binary classification task (at each instant of time, a loop closure is either detected or not), the cost of missed detections and false alarms are highly asymmetric: We pay a high price for declaring a loop closure that isn't, but there is minor harm in missing one, as temporal continuity affords many second chances in subsequent images. This is unlike large-scale image retrieval, where we wish to find what we are looking for (few missed detections, or high recall) even if we have to wade through some irrelevant hits (many false

alarms, or low precision).

Like image retrieval, however, the challenge with loop closure is scaling. In navigation applications, it may be hours before we return to a previously seen portion of the scene. Therefore, we have to store, and search through, hundreds of thousands to millions of images. Our goal in this paper is to *design a hierarchical data structure that helps speed up matching by leveraging on the two domain-specific constraints above:* temporal adjacency, and high precision.

Assuming continuous trajectories, the first translates to proximity in pose space $SE(3)$ (position and orientation). For the second, the best trade-off with missed detections can be achieved by testing every datum in the training set via *linear search accelerated via an inverse index*. Our goal is to achieve similar performance at a fraction of the cost compared to inverse index search. This cannot be achieved in a worst-case setting. What matters instead is *average performance* trading off precision with computational cost. We evaluate such average performance empirically on public benchmarks such as the *KITTI* dataset [11], the *City Centre*, and *New College* datasets [6], as well as demonstrating extensions to general image retrieval on *ukbench* [25] and *INRIA Holidays dataset* [16].

We propose a simple data structure based on hierarchical pooling of location likelihoods – in the form of sample distributions of visual descriptors – with respect to the topology of pose space. In practice, this means simply constructing bag-of-words of visual descriptors, that represent the likelihood of the locations that generated them, and pooling them temporally in a fine-to-coarse fashion, either by averaging, summing, or taking the index-wise maximum.

While averaging likelihoods may seem counter-productive, in Sec. 2 we show it makes sense in the context of the classical theories of sampling and anti-aliasing. In Sec. 3 we show that, despite its simplicity, it works as well as sophisticated agglomerative schemes at a fraction of the effort. While it may seem that something this simple that works must have been done before, we could not find trace of it in the literature, reviewed in Sec. 1.1, perhaps owing to the counter-intuitive nature of building hierarchies of pooled likelihoods.

1.1. Related work

Loop closure is a key component in robotic mapping (SLAM) [43], autonomous driving, location services on hand-held devices, and for wearables such as virtual reality displays. Appearance-only methods [7] are essentially large-scale retrieval algorithms, heavily influenced by [25] and more in general the literature of bag-of-word object recognition and categorization [31].

For scalability, the most common choice is to combine quantized local descriptors into a bag-of-words and then use an inverted index. FAB-MAP [6] extends the basic setup by learning a generative model of the visual words using a Chow-Liu tree to model the probability of co-occurrence of visual words. FAB-MAP 2.0 [7] scales further by exploiting sparsity to make the inverted index retrieval architecture more efficient. Angeli *et al*. [2] adopt Recursive Bayesian filtering with a simple voting scheme based on term frequency-inverted document frequency (tf-idf) [3, 1, 31]. Starting from [10], SIFT or SURF descriptors were replaced by more efficient binary descriptors such as BRIEF [4] and ORB [28] to achieve comparable precision and recall to FAB-MAP 2.0 with an order of magnitude speed increase. Several recent mapping and localization systems adopt it as a module, including [20] and ORB-SLAM [22].

In addition to the specific loop closure literature, general ideas from spatial data structures and agglomerative clustering [34] are also relevant to this work, including the KD-Tree [29], dual trees and decision trees [12], as well as data structures used for retrieval such as pyramid matching [14] and its spatial version [19]. In more general terms, this work also relates to visual navigation and mapping, structure-from-motion, and location recognition, including the use of global descriptors [37].

Our method can be considered appearance-only, but it is loosely informed by geometry, in the sense that the scene domain (pose space) provides the topology with respect to which we pool descriptors. Also closely related to our approach are [39, 36], which present techniques for merging only pairs of BoWs. To the best of our knowledge, our work is the first to pool them hierarchically and specifically target search time in loop closure.

2. Methodology

Since our focus is on a spatial structure that facilitates loop closure query, we adopt components from recent state-of-the-art methods as a baseline, and compare such a baseline against the same components integrated within our data structure. Specifically, we adopt [23] as a baseline, consisting of a bag-of-words where each word is an element of a dictionary of descriptors obtained off-line by hierarchical k-means clustering, with each word weighted by its inverse document frequency. FAST detectors [27] and BRIEF descriptors [4] are employed, and TF-IDF [3, 1, 31] is used to classify the BoW relative to the inverse index. This standard pipeline, with different clustering procedures to generate the dictionary and different features, comprises most basic large-scale retrieval systems, including appearance-only loop closure. However, the number of false alarms in large-scale settings is crippling, so temporal consistency and geometric verification are typically used as correction mechanisms. In a topological loop closure setting, matches against different images that are adjacent in time are lumped together as a single match, and highest-scoring matches are

checked for temporal consistency with previous time series data.

2.1. Hierarchical testing

Our data structure can be interpreted as a hierarchical version of TF-IDF. To illustrate the method, we first assume that every frame is a "keyframe" and therefore we have a time-series of BoWs, obtained as described above, and organized into a linear structure or *un-oriented list*, as we wish to retrieve regardless of the direction of traversal. Each node is associated with a histogram, in the form of a BoW, representing the likelihood of a pose $g(t) \in SE(3)$ (position $T(t) \in \mathbb{R}^3$ and orientation $R(t) \in SO(3)$) given the data (the image at time t, $I(t)$): $h^t \doteq \text{BoW}(t) \sim p(I(t)|g(t))$, where the equivalence is up to normalization, and the density function is approximated with a histogram with N bins, equal to the size of the dictionary.

We now construct a second level, or "layer", of the data structure, simply by pooling adjacent histograms (Fig. 1). This is repeated for higher layers until either a maximum depth is reached, or until a single root node is left. Similarities between pooled and query histograms are computed using the *intersection kernel* [33], that is the area of the intersection of the two histograms. Thus, if h^q (a query histogram) has bin values $h^q_1, \dots, h^q_N$, and similarly for h^p, we have that

$$\mathbb{I}(h^q, h^p) = \sum_{i=1}^{N} \min\{h^q_i, h^p_i\} \tag{1}$$

The intersection kernel is related to many divergence functions [41] as well as to metrics used in optimal transport problems.

Several standard choices for the pooling operation are available which we show trade off between precision and cost (Sec. 3). Suppose h^p is the parent histogram which has child histograms $\{h^k\}, k = 1, 2 \cdots K$. Both h^p and $h^k \in \mathbb{R}^N$. Mean- or average-pooling refers to $h^p = \frac{1}{K} \sum_{k=1}^{K} h^k$, sum-pooling refers to $h^p = \sum_{k=1}^{K} h^k$, and max-pooling refers to $h^p_i = \max_k \{h^k_i\}$, where $i = 1, 2 \cdots N$. The latter two pooling operators have an upper bound property when using the intersection kernel which makes the hierarchical test *have the same precision-recall behavior as standard inverse index search but also reasonable speedup*, though not as much speedup as the average operator. It is well known that for a query h^q,

$$\mathbb{I}(h^q, h^p) > \mathbb{I}(h^q, h^k) \tag{2}$$

when sum/max operators are used and thus if $\mathbb{I}(h^q, h^p) < \tau$, $\mathbb{I}(h^q, h^k) < \tau$ must hold. The downside of using sum/max-pooling is that without normalization, the pooled histograms may score high very quickly and trigger detections everywhere in the tree, which limits maximum speedup attainable.

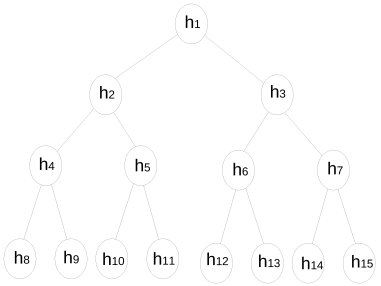

Figure 1: *Illustration of the construction of the hierarchical data structure for an 8-long sequence of (key)frames and constant branching factor of 2.*

Different trees with different depths and different branching factors can be constructed, trading off expected risk and computation time, characterized empirically in Sec. 3. In addition to a fixed depth (the data structure is then applied to the time series on a sliding window-based) and branching factor, one could devise more clever schemes to determine the topology of the tree, discussed in Sec. 2.3. However, we find that the benefit is limited compared to the straightforward fixed-topology architecture.

Suppose in our tree structure, a sliding window of K histograms $\{h^k\}, k = 1, 2 \cdots K$ are pooled into a parent node h^p, then h^p is an approximation of these histograms under one of the pooling operations. Based on the assumption that the child histograms are locally similar, the pooled histogram should be a good approximation to them. Since our goal is to search for the closest match, or at least for all matches that exceed a threshold $\tau > 0$ (we seek large values of $\mathbb{I}$), if $\mathbb{I}(h^q, h^p) < \tau$, the chance of any of h^p's descendants exceeding the threshold is rare (or impossible, in the case of max- or sum-pooling), therefore we stop searching the sub-tree rooted at h^p.

Therefore, search in a hierarchical TF-IDF setting simply boils down to *greedy breadth-first search, while maintaining an inverted index for each layer.* If only one layer is used, this reduces to standard linear search using an inverted index.

2.2. Keyframes and subsampling considerations

So far we have assumed that the time-series of data $\{h^t\}_{t=1}^{T}$ is sampled regularly (at constant time-intervals, or constant space-intervals), but it can also be sampled adaptively, by exploiting statistics of the data stream to decide which samples, or *keyframes*, to use. The data structure above does not change, since all that is required is a topology or adjacency structure to construct the tree. In Sec. 2.3 we explore more sophisticated heuristics to this end.

Adaptive (sub)-sampling can be done in a number of ways. The most popular involve a combination of heuristics that includes the average number of tracked features and

their overlap (new keyframes are spawned when the overlap of tracks with neighboring frames drop below a threshold), as well as hard gates on minimum and maximum time intervals, minimum and maximum distance in space or pose, between keyframes, including keyframes already existing in the map (but not temporally adjacent).

While there is significant performance or computational cost modulation to be had, our goal here is not to determine the best method for selecting keyframes. Instead, we focus on the data structure regardless of the sub-sampling mechanism. Consequently, we limit ourselves to constructing it either on the raw time series, or on any subsampling of it, as generated by any keyframe selection method.

It should be noted that mean-pooling, as described in the previous section, appears to be counter-productive since our goal is to maintain high precision. Just like selecting keyframes, building the hierarchy can be understood as a form of (sub)-sampling. Regardless of whether subsampling is regular (as in building the tree above) or adaptive (as in selecting keyframes), classical sampling theory [32] suggests that what should be stored at the samples is *not* the value of the function, but the local average relative to the topology of the domain where the data are defined (*anti-aliasing*).

In our case, the domain is time, or the order of keyframes, as a proxy of location in $SE(3)$. The range of the data is the space of likelihood functions, approximated by the BoW histograms h^t. Therefore, anti-aliasing simply reduces to averaging neighboring BoWs. The study of the optimal averaging, both in terms of support and weights, is beyond our scope here, where for mean-pooling we simply average nearest neighbors in the tree topology relative to a uniform prior. We do not delve into considering more sophisticated anti-aliasing schemes, since we have found that simple topologies yield attractive precision-computational cost tradeoff, which is unlikely to be significantly disrupted by fine-tuning the weights.

The practice of averaging likelihood functions as a way of anti-aliasing descriptors has also been recently shown by [8] in the context of pooling local descriptors for correspondences in wide-baseline matching. Our method can be considered an extension (or special case) where the correspondence and pooling are performed in time, and the descriptors are histograms of visual words, a mid-level representation, rather than histograms of gradient orientation, the result of low-level processing.

While our method is agnostic to the choice of heuristic to select keyframes, and in Sec. 3 we report results both with and without keyframes, it should be mentioned that even in the case of regular sampling, often for computational reasons datasets come down-sampled to the point where there is little to no overlap of descriptors or tracks between adjacent frames, that can therefore be safely considered all keyframes. One example (Oxford) will be illustrated in Sec. 3, where frames are so distant that tracking is not possible, and therefore wide-baseline correspondence must be employed for data association.

2.3. Adaptive tree topology

While the choice of heuristics for keyframe selection has no effect on our method, which can be applied to the raw time series or to the sequence of keyframes, the same (adaptive sampling) heuristics used to (down)-sample keyframes from the regularly sampled images could be used to aggregate nodes at one level into parents one level above. This would give rise to trees having different levels of connectivity at different layers, and indeed potentially at each node.

We have found that, in practice, these heuristics fail to yield significant performance improvements when compared to trees with fixed topology having constant splitting factors that match the average of their adaptive counter-part. Representative experiments are shown in Sec. 3, with more details in the Supplemental Material.

2.4. Generalization to other tasks

Although our approach is geared towards the loop closure scenario, its usage is not be restricted to it. A hierarchical structure of this form could be built on top of any histogram-based representation of images where some proxy of topology is available. In more general settings when a temporal stream of images is unavailable, extra labeling information, such as geotags, class labels, or textual annotations could be used. A hierarchy can be constructed using affinity between these alternate forms of metadata, provided that affinity implies proximity in the solution space. In Sec. 3.9, we demonstrate our data structure on two typical image retrieval benchmarks.

3. Evaluation

3.1. Datasets for loop closure

A common loop closure benchmark is the *KITTI* odometry dataset, consisting of 22 sequences of stereo pairs captured from a car driving through urban scenes [11].

We also test our scheme on the *Oxford City Centre* and *Oxford New College* datasets, captured by a hand-held camera [6]. These two datasets are challenging due to the ubiquity of repetitive structures such as brick walls and long corridors.

The datasets have different temporal sampling, finer in *KITTI* and coarser in *Oxford*. On the latter, each frame is considered a keyframe since the diversity between frames is such that each passes all the thresholds imposed by the most common keyframe selection heuristics.

3.2. Baseline algorithm

We use components of the recent ORB-SLAM [22], made available by the authors, as the baseline for our experiments. ORB-SLAM incorporates some recent techniques from binary feature descriptors and graph optimization. Modules in ORB-SLAM include FAST keypoint detection [35], ORB descriptors [28] and conversion to a bag-of-words. In loop detection, to quickly retrieve bag-of-words of previously visited places, an inverted index structure is maintained. ORB-SLAM also provides a heuristic to select keyframes, which we also use in our keyframe based experiments.

3.3. Choice of threshold

Missed detection and false alarms are traded off by an arbitrary choice of threshold, as in any detection algorithm.

Since the threshold affects the average query time (we can make that quite short by choosing a threshold that yields no false alarms while rejecting every hypothesis) we must come to a reasonable choice. We adopt the following policy: We generate precision-recall curves for all possible values of the threshold on *KITTI* 00. Then, we select the smallest threshold that yields zero false alarms (of course, that will still generally yield a non-zero false alarm rate in datasets that are not used in setting the threshold) and use it on other sequences.

We stress our focus on minimizing search-time of the query, and the data structure that enables it, rather than detection performance which is fixed by the choice of algorithm. Choosing sample images as queries ensures that the chosen algorithm will positively detect (at least) a loop closure when one is there. However, to avoid criticalities in the evaluation, we remove the query image from the dataset, but assume that there is a true match with neighboring images or keyframes. In some cases, the neighboring keyframes may be sufficiently different to fail the test, hence introducing spurious errors in the evaluation as these should be considered loop closures. This occurs, for instance, when the dataset is severely temporally down-sampled, as in the Oxford set. However, this handicap is shared by all methods, thus making our evaluation of the relative improvement due to the adoption of our data structure, fair.

3.4. Tradeoff between performance and cost

We show that our algorithm can be up to one order of magnitude faster than the baseline algorithm while still maintaining comparable performance in terms of precision and recall. The datasets we use are the left stereo images from *KITTI* sequences 00 and 02, which include 6-DoF ground truth poses and many loop closures. The evaluation protocol is as follows: traverse the sequence and insert BoW of images into database incrementally, while using each image to query the database. Two images are regarded as a correct match if they were taken within 15 meters of each other. To avoid trivial matches, we prevent the query from matching temporally adjacent images including the query itself. Fig. 2a and Fig. 2b show precision-recall curves of baseline algorithm and our hierarchical search with different topologies. Fig. 3 show timings of the baseline and our algorithm with different topologies and pooling schemes at the same threshold, chosen to achieve high precision and recall. Only time to query the database is counted, time for feature extraction and conversion to BoW are excluded. We use $d_i b_j$ -X to denote a hierarchy with i layers, a branching factor of j and pooling strategy X. Note that for baseline and our proposed algorithm with configuration $d_2 b_4$ -mean, $d_2 b_8$ -mean, the precision-recall curves are almost identical, while our approach is 2-5 times faster. For configuration $b_2 d_{16}$ -mean, while its performance is slightly worse, it achieves an order of magnitude speedup relative to the baseline. As mentioned in Sec. 2.1, sum/max-pooling have the same precision-recall behavior as the baseline. In these two datasets, sum- and max-pooling are slightly slower than inverted index search. Since both of these operations rapidly reduce sparsity in the histograms, we expect slower performance relative to mean-pooling. However, experiments in Sec. 3.7 and Sec. 3.9 show that sum/max-poolings do still speedup search in general.

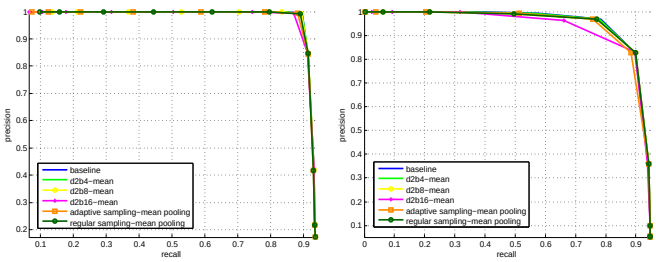

(a) *KITTI* 00 (b) *KITTI* 02

Figure 2: *Precision-recall curves of baseline and proposed algorithm with different topologies on* KITTI *dataset 00 and 02.* $d_i b_j$ *-X: a hierarchy with* i *layers, a branching factor* j *and pooling strategy X. Adaptive sampling: spectral clustering in* $SE(3)$ *with mean pooling. Regular sampling: sampling at the average rate of adaptive sampling scheme with mean pooling. Baseline: inverted index search.*

3.5. Adaptive domain-based clustering

In addition to the baseline algorithm, we generate a second baseline by applying the same algorithm to keyframes, rather than to all stored images. In principle, the heuristics involved in the selection of keyframes could be propagated to all nodes of the data structure, as discussed in Sec. 2.3. However, our experiments indicate that this yields minor benefits compared to simple averaging. The second

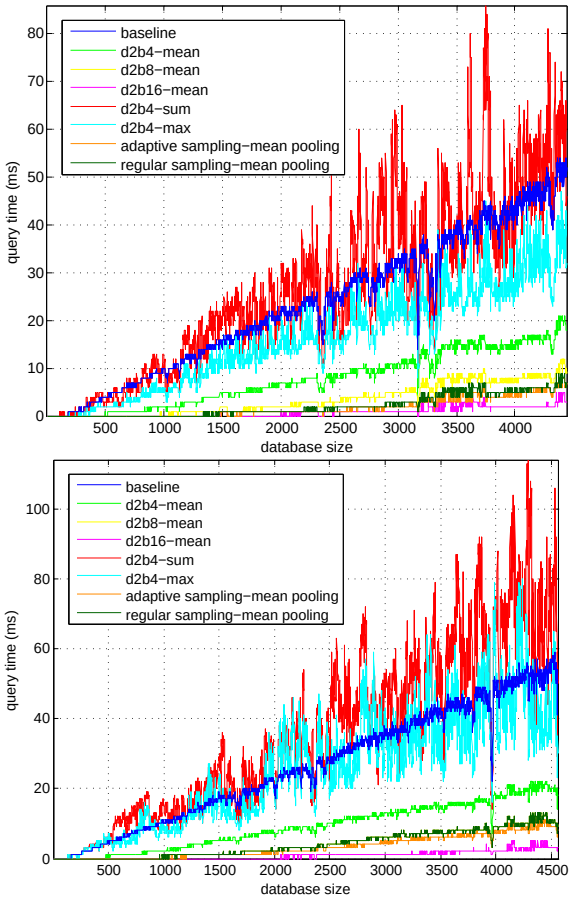

Figure 3: *Timings of baseline and proposed algorithm with different topologies and pooling strategies on* KITTI *dataset 00 and 02. Top:* KITTI *00. Bottom:* KITTI *02. Notations have the same meanings as in Fig. 2*

row of Table 2a shows average time-cost rate [1] for searching via an inverted index among keyframes, which is worse than searching in a simple hierarchy built on raw images, as shown in the second row of Table 1c. Besides, a simple regular sampling strategy on top of keyframes can speedup searching by a large margin as shown in Table 2a. Instead of a fixed topology of the data structure, corresponding to regular grouping, we can consider adaptive grouping based on a variety of criteria. Adaptive sampling, or grouping, based on *geometry* includes performing spectral clustering in $SE(3)$. Curves in Fig. 3 indicate that adaptive sampling achieves marginal improvements compared to regular sampling at a constant rate equal to the average of the adaptive sampling rate. Similarly, parallax-based sampling, based on clustering only the translational component of pose, also yields underwhelming improvements.

[1] Time-cost rate is defined as the increase of query time per thousand (1k) images in the database. Average time-cost rate is the average of time-cost rates computed for all sequences in each dataset.

3.6. Adaptive range-based clustering

Instead of clustering descriptors based on the domain where they are defined in a loop-closure system, that is pose space and the associated visibility and parallax, we can cluster descriptors based on where they stand in descriptor space. This is typically done in a discriminative setting using Agglomerative Information Bottleneck (AIB) [42, 34], which can also be used to reduce the complexity of the dictionary we use to compose the bag-of-words [9]. Unfortunately, the scale of operations desired prevents straightforward application of agglomerative clustering such as AIB, as well as standard data structures such as the KD-Tree or decision trees. Also, dimensionality reduction by clustering in descriptor space is counter-productive, because the goal is to discriminate pose, and therefore similar descriptors that refer to distant points in pose space should not be averaged (think of driving through a countryside where the landscape changes slightly over long stretches of open road), whereas different descriptors generated by nearby locations (think of rotating in place in a cluttered indoor scene) can be safely aggregated because loop closure can be called so long as any of them triggers a match.

3.7. Quantifying speedup using synthetic ground-truth

Depending on the particular query image, our method could reduce or increase search time. The former occurs when correspondence fails early allowing us to rule out subsequent tests at finer scales. However, in the worst-case we may end up performing more comparisons than inverse index search when the test reaches the finest scale too often. In practice, what matters is that our algorithm shortens test time *on average* during long sequences. Since most *KITTI* sequences contain few or no loop closures, we generate synthetic positive and negative queries as follows: For sequences 01 to 21, we generate positive queries by sampling the right stereo images of each sequence (slightly different from the left images from which we constructed the database), and generate negative queries by sampling images from sequence 00. For the *Oxford* datasets, we construct the database using odd-numbered images, generate positive queries from the even-numbered images, and negative queries again from *KITTI* 00.

Overall performance is measured by combining both sets of queries. Of course, even in the negative case our algorithm could find erroneous correspondences, which are then labeled as false alarms. Similarly, we may find no correspondence in the former case (missed detection). We use average time-cost rate to evaluate how the searching algorithm scales with size of the database.

Table 1a and 1b show average time-cost rate and speedup over inverse index search on the *KITTI* dataset for positive/negative queries respectively. Overall performance,

structure	pooling	rate(ms/1k)	speedup
inverted index	N/A	10.07	1.00
hierarchical	mean	**0.69**	**14.59**
	sum	8.70	1.16
	max	6.65	1.52

(a) performance when only positive queries are used

structure	pooling	rate(ms/1k)	speedup
inverted index	N/A	9.86	1.00
hierarchical	mean	**0.34**	**29.00**
	sum	6.28	1.57
	max	5.04	1.96

(b) performance when only negative queries are used

structure	pooling	rate(ms/1k)	speedup
inverted index	N/A	9.88	1.00
hierarchical	mean	**0.38**	**26.00**
	sum	7.92	1.25
	max	6.06	1.63

(c) overall performance

Table 1: *Average time-cost rate and speedup over 21 sequences of* KITTI dataset. *1st col: grouping strategies. 2nd col: pooling operations. 3rd col: average time-cost rate, which describes how the query time increases per 1k images inserted into the database.*

structure	pooling	rate(ms/1k)	speedup
inverted index	N/A	8.97	1.00
hierarchical	mean	**0.88**	**10.14**
	sum	7.87	1.14
	max	6.00	1.50

(a) KITTI

structure	pooling	rate(ms/1k)	speedup
inverted index	N/A	6.98	1.00
hierarchical	mean	**1.61**	**4.34**
	sum	4.71	1.48
	max	4.20	1.66

(b) Oxford

Table 2: *A comparison of search in flat and hierarchical structure on* KITTI *and* Oxford *dataset. Notations have the same meanings as in Table 1 except that 3rd column describes average time-cost rate over the 21* KITTI keyframe *sequences and all 4 sequences in the* Oxford *datasets respectively. The keyframes are generated by running ORB-SLAM.*

combining positive and negative sets, is reported in Table 1c for *KITTI* and Table 2b for the *Oxford* datasets. Also, we repeat the same experimental procedure on a subset of the original database selected using heuristic keyframe selection, instead of using every image of the sequence, in Table 2a.

3.8. Empirical Assessment

In general, it seems that the domain-specific constraints would favor deep structures with small branching factors for the purpose of accelerating the query. Assuming that loop closures are rare, which is the case in long driving sequences, we can reject queries by just visiting a few top-level nodes of the hierarchy which is cheaper compared to visiting a lot of bottom-level nodes. However, this only holds for the hierarchy without inverted index. With the inverted index, to retrieve a bag-of-words is just to retrieve sparse common words shared by query and database entries, which makes visiting extra nodes very cheap. On the other hand visiting an extra layer of the hierarchy always imposes overhead, which might not be negligible depending on the dataset. Therefore, when an inverted index is employed, a shallow structure with suitable branching factor depending on the application is preferred. Furthermore, for short-term loop closure, as employed during tracking and mapping, the priorities are different. But since in that case scale is not as cogent an issue, we focus on large-scale, rare-loop scenarios. Fig. 4 shows performance of different tree topologies using mean-pooling on *KITTI* 02. All sampled topologies improve the baseline by a considerable margin. Depending on the application context and constraints, one may favor one architecture over another. This determination has to be made empirically. Here we have focused on benchmark datasets that are commonly employed for navigation, mapping, and loop closure applications.

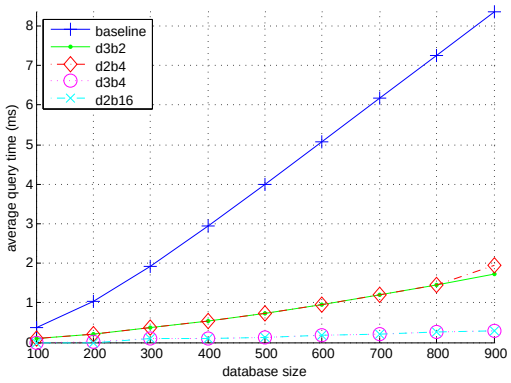

Figure 4: *Performance of different combinations of depth and branching factor using the experimental procedure of Sec. 3.7 on* KITTI *02. Baseline: basic inverted index; $d_i b_j$: tree structure with depth i and branching factor j. Mean-pooling is used.*

3.9. Experiments in image retrieval tasks

We test our algorithm on two publicly available image retrieval benchmarks: *ukbench* [25] and *INRIA Holidays dataset* [15].

ukbench [2] consists of 2550 groups of 4 images each (10200 total). The 4 images in each group contain the same object or location, but under different viewpoint, rotation, scale and lighting conditions. We use the same evaluation protocol available on the authors' website: count how many of 4 images are top-4 when using a query image from that set of four images. We use the pre-computed visual words provided by the authors, which are quantized SIFT descriptors using a 1M vocabulary.

[2] http://vis.uky.edu/~stewe/ukbench/

INRIA Holidays dataset [3] contains 500 image groups (1491 total), each of which represents a distinct scene or object under different rotations, viewpoint and illumination changes, blurring, etc. The dataset covers a very large variety of scene types (natural, man-made, water and fire effects, etc). Performance is measured by mean average precision (mAP) averaged over all 500 queries. In our experiments, we use the 4.5 million SIFT descriptors and 100k vocabulary provided by the authors to construct a BoW for each image.

The baseline remains search using an inverted index system. Our algorithm is facilitated by a three-layer data structure with original histograms at the bottom layer. At the second layer, histograms belonging to the same object/scene are pooled (pooling based on prior information available about the data and problem space). At the top layer, we compare two different strategies to build the hierarchy: random grouping and greedy affinity grouping. Random grouping: we randomly group every N histograms from the second layer. Greedy affinity grouping: we greedily group every N histograms based on their nearest neighbors in affinity (which is the histogram intersection score). In each setup, we also compare the different choices of pooling operators. Table 3a and Table 3b show results on the *ukbench* and *Holidays* datasets with $N = 16$, more experiment results with different hierarchical structures can be found in Supplemental Material.

The search strategy we use in image retrieval is a slightly different from what we use in loop closure. Unlike the hierarchical test in loop closure which we search down those nodes which score higher than the threshold, in image retrieval tasks, we completely discard the threshold and only search down those nodes which have top 10 highest scores. Thus even for sum/max-poolings, the precision-recall behavior should be different from the baseline. All hierarchical approaches, regardless of pooling operation, are faster than the baseline. We also notice improved score/mAP in these two experiments, likely due to the grouping of histograms of the same object/scene at the second layer of our hierarchy and the top-4 scoring mechanism imposed by the benchmark.

structure	pooling	time(ms)	speedup	score
inverted index	N/A	1.47	1.00	2.72
Random hierarchical	mean	0.38	3.87	2.80
	sum	**0.37**	**3.97**	**2.83**
	max	0.39	3.77	2.82
Greedy affinity hierarchical	mean	0.38	3.87	2.80
	sum	0.38	3.87	**2.83**
	max	**0.37**	**3.97**	2.82

(a) ukbench

structure	pooling	time(ms)	speedup	mAP
inverted index	N/A	9.11	1.00	0.56
Random hierarchical	mean	**5.57**	**1.63**	0.58
	sum	6.19	1.47	**0.63**
	max	6.24	1.46	0.62
Greedy affinity hierarchical	mean	5.58	1.63	0.57
	sum	6.82	1.34	**0.63**
	max	6.53	1.40	0.62

(b) Holidays dataset

Table 3: *A comparison of search in flat and hierarchical structure on* ukbench *and* Holidays dataset. *1st col: grouping strategies. 2nd col: pooling operations. 3rd col: average query time.* ukbench *takeks average number of top-4 retrieved images as score.* Holidays dataset *takes mAP as evaluation metric.*

4. Discussion

We have presented a hierarchical data structure consisting of pooled local descriptors representing the likelihood of locations given the images they generate, while maintaining an inverted index at each level of the data structure. While mean-pooling of histograms may seem counter-productive, it is a sensible choice when considered an anti-aliasing procedure in the context of classical sampling theory, where the data structure, as well as keyframes, are tasked to *down-sampling* the native rate. We have compared several pooling strategies, and found that mean-pooling provides most speedup at small cost to performance; sum-pooling has the upper-bound property and accelerates search to a reasonable degree without loss of performance; and max-pooling shares the same property with sum-pooling but exhibits a larger speedup due better approximating the nodes below it.

For simplicity, we chose a fixed topology (depth and branching factor) and studied the resulting performance empirically. We have found that sophisticated heuristics do not improve performance enough to justify the added complexity. We have benchmarked our scheme on public datasets, where we have shown that even a shallow tree can significantly cut down on test time with minimal impact to precision, which is the main goal of loop closure.

[3] https://lear.inrialpes.fr/~jegou/data.php

References

[1] A. Aizawa. An information-theoretic perspective of tf–idf measures. *Information Processing & Management*, 39(1):45–65, 2003. 2

[2] A. Angeli, D. Filliat, S. Doncieux, and J.-A. Meyer. Fast and incremental method for loop-closure detection using bags of visual words. *Robotics, IEEE Transactions on*, 24(5):1027–1037, 2008. 2

[3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003. 2

[4] M. Calonder, V. Lepetit, C. Strecha, and P. Fua. Brief: Binary robust independent elementary features. In *Computer Vision–ECCV 2010*, pages 778–792. Springer, 2010. 2

[5] O. Chum, J. Philbin, J. Sivic, M. Isard, and A. Zisserman. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007. 1

[6] M. Cummins and P. Newman. Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research*, 27(6):647–665, 2008. 1, 2, 4

[7] M. Cummins and P. Newman. Highly scalable appearance-only slam-fab-map 2.0. In *Robotics: Science and Systems*, volume 5. Seattle, USA, 2009. 2

[8] J. Dong and S. Soatto. Domain size pooling in local descriptors: Dsp-sift. In *Proc. IEEE Conf. on Comp. Vision and Pattern Recogn.*, 2015. 4

[9] B. Fulkerson, A. Vedaldi, and S. Soatto. Localizing objects with smart dictionaries. In *Proc. of the Eur. Conf. on Comp. Vis. (ECCV)*, October 2008. 6

[10] D. Galvez-Lopez and J. D. Tardos. Real-time loop detection with bags of binary words. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 51–58. IEEE, 2011. 1, 2

[11] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, page 0278364913491297, 2013. 2, 4

[12] D. Geman and B. Jedynak. An active testing model for tracking roads in satellite images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 18(1):1–14, 1996. 2

[13] R. Girshick, F. Iandola, T. Darrell, and J. Malik. Deformable part models are convolutional neural networks. *arXiv preprint arXiv:1409.5403*, 2014. 1

[14] K. Grauman and T. Darrell. The pyramid match kernel: Discriminative classification with sets of image features. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 2, pages 1458–1465. IEEE, 2005. 2

[15] H. Jegou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *Computer Vision–ECCV 2008*, pages 304–317. Springer, 2008. 7

[16] H. Jégou, M. Douze, and C. Schmid. Recent advances in large scale image search. In *Emerging Trends in Visual Computing*, pages 305–326. Springer, 2009. 1, 2

[17] H. Jégou, M. Douze, and C. Schmid. Improving bag-of-features for large scale image search. *International Journal of Computer Vision*, 87(3):316–336, 2010. 1

[18] E. Jones and S. Soatto. Visual-inertial navigation, localization and mapping: A scalable real-time large-scale approach. *Intl. J. of Robotics Res.*, april 2011. 1

[19] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 2169–2178. IEEE, 2006. 2

[20] H. Lim, J. Lim, and H. J. Kim. Real-time 6-dof monocular visual slam in a large-scale environment. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 1532–1539. IEEE, 2014. 2

[21] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *Proc. IEEE Conf. on Comp. Vision and Pattern Recogn.*, 2014. 1

[22] R. Mur-Artal, J. Montiel, and J. D. Tardos. Orb-slam: a versatile and accurate monocular slam system. *Robotics, IEEE Transactions on*, 31(5):1147–1163, October 2015. 1, 2, 5

[23] R. Mur-Artal and J. D. Tardós. Fast relocalisation and loop closing in keyframe-based slam. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 846–853. IEEE, 2014. 2

[24] R. A. Newcombe and A. J. Davison. Live dense reconstruction with a single moving camera. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1498–1505. IEEE, 2010. 1

[25] D. Nister and H. Stewenius. Scalable recognition with a vocabulary tree. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 2161–2168. IEEE, 2006. 1, 2, 7

[26] A. B. Patel, T. Nguyen, and R. G. Baraniuk. A probabilistic theory of deep learning. *arXiv preprint arXiv:1504.00641*, 2015. 1

[27] E. Rosten and T. Drummond. Machine learning for high-speed corner detection. In *Computer Vision–ECCV 2006*, pages 430–443. Springer, 2006. 2

[28] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: an efficient alternative to sift or surf. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2564–2571. IEEE, 2011. 2, 5

[29] H. Samet. *The design and analysis of spatial data structures*, volume 85. Addison-Wesley Reading, MA, 1990. 2

[30] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *Workshop at International Conference on Learning Representations*, 2014. 1

[31] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 1470–1477. IEEE, 2003. 2

[32] S. Smale and D.-X. Zhou. Shannon sampling ii: Connections to learning theory. *Applied and Computational Harmonic Analysis*, 19(3):285–302, 2005. 4

[33] M. J. Swain and D. H. Ballard. Color indexing. *International journal of computer vision*, 7(1):11–32, 1991. 3

[34] N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. In *Proc. of the Allerton Conf.*, 2000. 2, 6

[35] E. Tola, V. Lepetit, and P. Fua. A fast local descriptor for dense matching. In *Proc. CVPR*. Citeseer, 2008. 5

[36] A. Torii, J. Sivic, and T. Pajdla. Visual localization by linear combination of image descriptors. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 102–109. IEEE, 2011. 2

[37] A. Torralba, K. Murphy, W. Freeman, and M. Rubin. Context-based vision system for place and object recognition. In *Proc. 9^{th} Int. Conf. on Computer Vision*, 2003. 2

[38] A. Torralba, K. P. Murphy, W. T. Freeman, and M. A. Rubin. Context-based vision system for place and object recognition. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 273–280. IEEE, 2003. 1

[39] P. Turcot and D. G. Lowe. Better matching with fewer features: The selection of useful features in large database recognition problems. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pages 2109–2116. IEEE, 2009. 2

[40] I. Ulrich and I. Nourbakhsh. Appearance-based place recognition for topological localization. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 2, pages 1023–1029. Ieee, 2000. 1

[41] N. Vasconcelos. On the efficient evaluation of probabilistic similarity functions for image retrieval. *Information Theory, IEEE Transactions on*, 50(7):1482–1496, 2004. 3

[42] A. Vedaldi and B. Fulkerson. Vlfeat: An open and portable library of computer vision algorithms. In *Proceedings of the international conference on Multimedia*, pages 1469–1472. ACM, 2010. 6

[43] B. Williams, M. Cummins, J. Neira, P. Newman, I. Reid, and J. Tardós. A comparison of loop closing techniques in monocular slam. *Robotics and Autonomous Systems*, 2009. 2